Improved Oracles for Time-Dependent Road Networks*

Spyros Kontogiannis[1,4], Georgia Papastavrou[1,4], Andreas Paraskevopoulos[3,4], Dorothea Wagner[2], and Christos Zaroliagis[3,4]

1 **Department of Comp. Science & Engineering, University of Ioannina, Greece** [kontog@cse.uoi.gr , gioulycs@gmail.com]

2 **Karlsruhe Institute of Technology, Germany** [dorothea.wagner@kit.edu]

3 **Department of Comp. Eng. & Informatics, University of Patras, Greece** [{ paraskevop , zaro }@ceid.upatras.gr]

4 **Computer Technology Institute and Press "Diophantus", Greece**

arXiv:1704.08445v1 [cs.DS] 27 Apr 2017

Abstract

A novel landmark-based oracle (CFLAT) is presented, which provides earliest-arrival-time route plans in time-dependent road networks. To our knowledge, this is the first oracle that preprocesses combinatorial structures (collections of time-stamped min-travel-time-path trees) rather than travel-time functions. The preprocessed data structure is exploited by a new query algorithm (CFCA) which computes (and pays for it), apart from earliest-arrival-time estimations, the *actual connecting path* that preserves the theoretical approximation guarantees. To make it practical and tackle the main burden of landmark-based oracles (the large preprocessing requirements), CFLAT is extensively engineered. A thorough experimental evaluation on two real-world benchmark instances shows that CFLAT achieves a significant improvement on preprocessing, approximation guarantees and query-times, in comparison to previous landmark-based oracles, whose query algorithms do *not* account for the path construction. It also achieves competitive query-time performance and approximation guarantees compared to state-of-art speedup heuristics for time-dependent road networks, whose query-times in most cases do *not* account for path construction.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems;

Keywords and phrases Time-dependent shortest paths; FIFO property; Distance oracles.

Digital Object Identifier 10.4230/LIPIcs...

1 Introduction

The surge for efficient solutions (min-cost paths) in networks with temporal characteristics is a highly challenging research goal, due to both the large-scale and the time-varying nature of the underlying arc-cost metric. Along this line, the development of practical algorithms for providing earliest-arrival-time route plans in large-scale road networks accompanied with a time-dependent arc-travel-time metric (known as *Time-Dependent Route Planning* – TDRP), has received a lot of attention in the last decade. TDRP is a hard challenge, both theoretically and in practice. For certain tractable cases, there is an analogue of Dijkstra's algorithm (called *Time-Dependent Dijkstra* – TDD) to solve the problem in quasi-linear time, which is already too much for a route-planning application supporting real-time query responses in *large-scale* road networks. Time-dependence is also by itself a

* Partially supported by EU FP7/2007-2013 under grant agreements no. 609026 (project MOVESMART), no. 621133 (project HoPE), and by DFG grant WA 654/23-1 within FOR 2083.

© S. Kontogiannis, G. Papastavrou, A. Paraskevopoulos, D. Wagner and C. Zaroliagis;
licensed under Creative Commons License CC-BY

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

quite important degree of complexity, both in space and in query-time requirements. These two challenges have been tackled in the past either by oracles, or by speedup heuristics. An *oracle* is a preprocessed and succinctly stored data structure encoding min-cost path information for carefully selected pairs of vertices. This data structure is accompanied with a query algorithm, which responds to arbitrary queries in time *provably better* than the corresponding Dijkstra-time and, if approximate solutions are also an option, with a *provable* approximation guarantee (stretch). Analogously, a *speedup heuristic* preprocesses arc-cost metrics which are custom-tailored to road networks, and then uses a query algorithm for responding to (exact or approximate) min-cost path queries in time that is in practice *several orders of magnitude* faster than the running time of Dijkstra's algorithm.

Modeling Instances, Problem Statement & Related Work. We model road network instances by directed graphs in which every arc $a = uv$ depicts an uninterrupted portion of a road segment and is accompanied by an arc-travel-time *function* $D[a]$ determining the time to traverse a, given the departure-time from its tail u. These functions are assumed to be continuous, piecewise-linear (pwl), periodic with one-day period, and are succinctly represented as sequences of consecutive *breakpoints*, i.e., (departure-time,arc-travel-time) pairs. This model is typical in the literature when we seek for route plans for private cars (e.g., [9, 10, 6, 19, 7, 3, 12, 20, 17, 16, 14, 4, 15]). For an arbitrary pair (o, d) of origin-destination points, there are two main algorithmic challenges: (i) $TDRP(o, d, t_o)$ concerns the computation of a minimum travel-time *od*-path for a given departure-time t_o, i.e., the *evaluation* of the minimum-travel-time function $D[o, d](t_o)$ from o to d; (ii) $TDRP(o, d)$ concerns the construction and succinct representation of the entire function $D[o, d]$, for *all* possible departure-times (e.g., for future instantaneous evaluations). A crucial property that makes $TDRP(o, d, t_o)$ tractable is the FIFO property, according to which delaying the departure-time from the tail of an arc cannot possibly cause an earlier arrival at its head (i.e., the arcs behave as FIFO queues). For FIFO-abiding instances, a time-dependent variant of Dijkstra's algorithm (`TDD`) running in quasi-linear time is known [11, 21]. Without the FIFO property the problem can become extremely hard, depending on the adopted waiting policy at the vertices of the network [21]. As for $TDRP(o, d)$, this is known to be hard even when the FIFO property holds [12]. Fortunately, if (good) upper-approximations $\overline{\Delta}[o, d]$ of the minimum-travel-time functions $D[o, d]$ are an option, then there exist polynomial-time and space-efficient *one-to-one* [5, 12, 20], or *one-to-all* [15, 16, 17] approximation algorithms.

As a quality measure, independent of the query at hand, the *relative error* is typically used, i.e., the *maximum absolute error* (MAE) divided by the optimal travel-time; the MAE is the worst-case difference of an optimal travel-time from the proposed (path's) travel-time.

Several speedup heuristics, with remarkable success in road networks possessing scalar arc-cost metrics, have been extended to the case of TDRP. Some of them [7, 8, 19] are based on (scalar) lower bounds of travel-time functions (e.g., free-flow travel-times) to orient the search for a good route. `TDCALT` [7] yields reasonable query-response times for $TDRP(o, d, t_o)$, and `TDSHARC` [6] provides in reasonable time solutions to $TDRP(o, d)$, even for continental-size networks. `TDCRP` [4] is currently one of the most successful speedup heuristics, whose main feature is *customizability*, i.e., almost real-time adaptation to changes in the arc-cost metric. `TCH` [3] also achieves remarkable query times, both for $TDRP(o, d, t_o)$ and for $TDRP(o, d)$, even for continental-size networks. All the above mentioned heuristics only compute (estimations of) erarliest-arrival-times, excluding the overhead for constructing the corresponding connecting path. The only heuristics that also account the path construction in their query-times are provided in [22], with quite competitive performances.

In parallel to speedup heuristics, there has been a recent trend to provide oracles for

TDRP, with *provable* theoretical performance and approximation guarantees [16, 17], which have been experimentally evaluated on real-world instances [14, 15]. The most successful one, `FLAT` [15, 16], demonstrated in practice noticeable query times and relative errors, much better than the theoretical guarantees, thus being competitive to the aforementioned speedup heuristics, justifying further research on providing even better oracles for TDRP, for the additional reason that oracles also ensure scalability.

Contributions and Outline. We present, engineer and experimentally evaluate `CFLAT` (Section 2), a novel *landmark-based* oracle for TDRP whose objective is to tackle the main burden of such oracles, the large preprocessing requirements, without compromising either the preprocessing scalability, the competitiveness of query-response times, or the approximation guarantees. To our knowledge, `CFLAT` is the first oracle for time-dependent networks that preprocesses only time-evolving *combinatorial structures*: it maintains a carefully selected collection of time-stamped min-cost-path trees which can assure good approximation guarantees while minimizing the required space. Computing (and storing) less during preprocessing, unavoidably leads to more demanding work per query in real-time. Nevertheless, our novel query algorithm (`CFCA`) manages to achieve better query times and significantly improved practical performance compared to previous oracles, despite the fact that it actually computes a connecting path, and not just an estimation of a good upper bound on the minimum travel-time for the query at hand, as is done by almost all other oracles and speedup techniques for TDRP. Our specific contributions are threefold: **(i)** We propose `CTRAP` (Section 2.2.1), a novel *approximation method* which stores only min-cost-path trees for carefully selected landmark vertices and sampled departure-times. Apart from the obvious economy of space due to omitting certain attributes (travel-time values), the novelty of this approach is that it exploits the fact that there are significantly fewer changes in the combinatorial structure, than in the functional description of the optimal solution (earliest arrival-times at a destination). Moreover, we avoid multiple copies of the same preprocessed information, by organizing the destinations from a landmark into groups of (roughly) equidistant vertices, for which the common departure-times sequence is stored only once. We then proceed with the *landmark selection policies* (Section 3) considered by `CFLAT`. Apart from the most successful ones in [15], we also consider new policies based on the betweeness-centrality measure. Due to the significant reduction in space requirements, we are in a position to select much larger landmark sets, which allows us to showcase the full scalability of `CFLAT` in trading smoothly preprocessing requirements with query response times and approximation guarantees. **(ii)** We propose `CFCA(N)` (Section 2.2.2), a novel *query algorithm* that exploits the preprocessed information of `CFLAT`: For a query (o, d, t_o), it starts by growing a `TDD` ball from o at time t_o, until the N closest landmarks are settled. It then marks a small subset of relevant arcs, using the N settled landmarks as "attractors" that orient the discovery of certain paths from d back to o. This is reminiscent of the `ARCFLAGS` algorithm for static metrics [13], but the choice of the relevant arcs is done "on the fly", since this information is also time-dependent. In the final step, it continues growing the initial `TDD` ball, but only within the subgraph of marked arcs, until the destination d is settled within this subgraph. `CFCA(N)` achieves the same theoretical approximation guarantee with the query algorithm `FCA(N)` of `FLAT`, but in practice it is much better than `FCA(N)`. **(iii)** We conduct a thorough *experimental evaluation* of `CFLAT` (Section 3), on two well established real-world instances, the urban area of Berlin and the national road network of Germany. Our findings are perceptible. For Berlin, the preprocessing requirements are less than 3.306sec and 2.521MB (0.69MB compressed) per landmark. Thus, if space is our primary concern, we can preprocess 250 random landmarks in less than 14min, consuming 0.7GB (0.17GB compressed) space, whereas the query per-

formance (average query time and relative error) varies from 0.565msec and 2.418% (for $N = 1$), to 3.330msec and 0.136% (for $N = 6$). With 16K landmarks the query performance varies from 0.076msec and 0.192% (for $N = 1$), to 0.226msec and 0.022% (for $N = 6$). As for Germany, the preprocessing requirements are 29.322sec and 26.8MB (8.07MB compressed) per landmark. For $4K$ landmarks, we achieve a query performance varying from 0.683msec and 0.831% (for $N = 1$), to 4.104msec and 0.031% (for $N = 6$).

2 The CFLAT Oracle

A landmark-based oracle selects a set $L \subseteq V$ of *landmarks* and preprocesses travel-time information (*summaries*) between them and all (or some) reachable destinations. A query algorithm exploits these summaries for responding to earliest-arrival-time queries (o, d, t_o), from an origin o and departure-time t_o to a destination d, in time that is *provably* efficient (e.g., sublinear in the size of the instance). The oracle is also accompanied with a *theoretically proved* approximation guarantee (a.k.a. stretch) for the quality of the recommended routes.

In Section 2.2 we present our novel oracle, CFLAT. Before doing that, we recap in Section 2.1 FLAT, an oracle upon which CFLAT builds and achieves remarkable improvements.

2.1 Recap of FLAT

FLAT is, to date, the most successful oracle for TDRP in road networks, and was originally presented and analyzed in [16]. A variant of FLAT was implemented and experimentally evaluated in [15]. In this work, we consider (and refer to as FLAT) to that variant. Its main building block is the TRAP approximation method: Given a landmark ℓ, the period $[0, T)$ is split into intervals of an (arbitrarily chosen) length $3,200$sec. The endpoints of these intervals are used as sampled departure-times. The corresponding min-cost-path trees rooted at ℓ are computed, producing travel-time values for all reachable destinations v. For each interval $[t_s, t_f)$, an upper-approximating function $\overline{\delta}$ is considered, which is the lower-envelope of a line of max slope ($\Lambda_{\max}$) passing via $\langle t_s, D[\ell, v](t_s)\rangle$ and a line of min slope ($-\Lambda_{\min}$) passing via $\langle t_f, D[\ell, v](t_f)\rangle$ (cf. Figure 1). Observe that $\overline{\delta}$ considers an *intermediate breakpoint* $\langle \overline{t}_m, \overline{D}_m\rangle$, the intersection of the two lines, which is *not* the outcome of an actual sampling. This intermediate breakpoint is only stored when v becomes deactivated (i.e., within this interval there is no need for further sample points, see next paragraph). A similar lower-approximating function $\underline{\delta}$ is considered, which is the upper-envelope of a min-slope line passing via $\langle t_s, D[\ell, v](t_s)\rangle$ and a max-slope line passing via $\langle t_f, D[\ell, v](t_f)\rangle$.

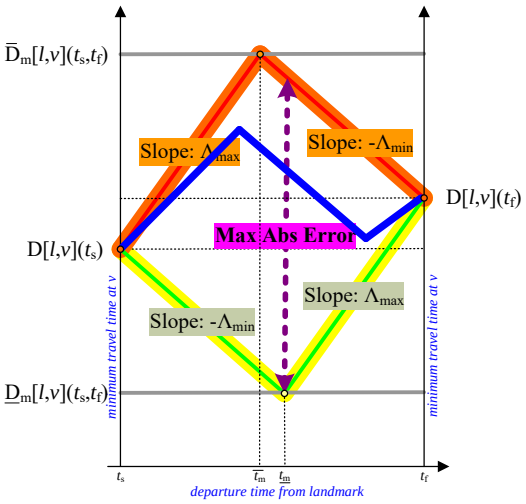

Figure 1 Upper-approximation $\overline{\delta}_k[\ell, v]$ (thick-orange) and lower-approximation $\underline{\delta}_k[\ell, v]$ (thick-green) of $D[\ell, v]$ (blue), within $[t_s, t_f)$.

A closed-form expression of the worst-case error (*maximum absolute error* – MAE) is used to determine whether $\overline{\delta}$ is a sufficient upper-approximation of $D[\ell, v]$ within $[t_s, t_f)$, given a required approximation guarantee $\varepsilon > 0$. If this is the case, v becomes *deactivated* for this subinterval, meaning that no more sampled trees will be of interest for v within it. TRAP continues by choosing finer sampling intervals, first of length $1,600$sec, then 800sec,

400sec, etc., computing min-cost-path trees only for the new departure-time samples in each round, until eventually there is no active destination for any of subintervals of the currently chosen length. The concatenation of all the upper-approximations for the smallest active subintervals of v is considered by TRAP as the required $(1+\varepsilon)$-upper-approximation $\overline{\Delta}[\ell, v]$ (called a *travel-time summary*) of $D[\ell, v]$ within $[0, T)$. $\overline{\Delta}[\ell, v]$ is stored as a sequence of pairs of breakpoints, i.e., (departure-time,travel-time) pairs, in increasing order w.r.t. departure-times. During the preprocessing, FLAT calls TRAP to produce travel-time summaries, from a carefully selected set of landmark vertices towards all reachable destinations.

Upon a query (o, d, t_o) FLAT calls FCA(N)[1], a query algorithm which grows a TDD ball from o with departure-time t_o, until either d or the first N landmarks are settled. It then returns either the exact route (when d is settled), or the best-of-N (w.r.t. the *theoretical guarantees*) od-path passing via one of the N settled landmarks and being completed (from ℓ to d) by exploiting the preprocessed summaries for d. Since FCA(N) does not need all summaries to be concurrently available in memory, the preprocessed data blocks representing travel-time summaries of FLAT were compressed, and only summaries of the landmarks required per query were decompressed on the fly. The zlib library was used for this purpose, leading to a reduction of 10% in the required space. More details on FLAT are provided in [15, 16].

2.2 Description of CFLAT

We now present CFLAT, which can be considered as the combinatorial analogue of FLAT. At a high level, CFLAT works as follows. In a preprocessing phase, it constructs and compactly stores min-cost-path trees at carefully sampled departure-times, rooted at each landmark $\ell \in L$. A query (o, d, t_o) is answered by first growing a TDD ball from o at time t_o, until either d or a small number of landmarks are settled. In the latter case, starting from d, a suitably small subgraph is constructed (consisting of certain paths going from d back to o, using the settled landmarks as "attractors"), until a settled vertex of the initial TDD ball is reached. Then, a continuation of growing the initial TDD ball on the resulted small subgraph returns an od path that turns out to approximate very well the optimal od path.

2.2.1 The Approximation Method CTRAP and CFLAT Preprocessing

CTRAP computes and stores only min-cost-path trees at carefully sampled departure-times, rather than actual breakpoints of the corresponding minimum-travel-time functions. The algorithm's pseudocode is provided in the appendix (cf. Section B). We present here only a sketch of the main steps as well as the key new insights, compared to TRAP. CFLAT preprocessing consists simply in calling CTRAP(ℓ, ε) for each landmark $\ell \in L$.

procedure CTRAP(ℓ, ε)

STEP 1: Keep sampling finer departure-times from $[0, T)$, as in TRAP, until all destinations achieve relative error less than ε and become inactive.
1.1: Store (pruned at inactive nodes) min-cost-path trees from ℓ, for all departure-times.
1.2: Omit intermediate breakpoints.

STEP 2: Merge consecutive breakpoints with identical predecessors.

STEP 3: Avoid multiple copies of common departure-time sequences.

When executed from a landmark ℓ, CTRAP works as follows: Step 1 resembles TRAP, the only difference being that CTRAP keeps only the immediate predecessors (parents) per active destination v in the sampled min-cost-path trees. In particular, a pair of sequences is created, $PRED[\ell, v]$ for predecessors and $DEP[\ell, v]$ for the corresponding sampled departure-times,

[1] In [15] it was called FCA$^+$, with a fixed number $N = 6$ of landmarks to settle.

per landmark-destination pair $(\ell, v) \in L \times V$. Step 2 cleans up each pair of sequences, by merging consecutive breakpoints for which the predecessor is the same. Step 3 organizes the destinations from a landmark ℓ into groups with the same departure-times sequence, so that multiple copies of the same sequence are avoided. In the rest of this section, we describe in more detail the key new insights and algorithmic steps of CTRAP, compared to TRAP [15, 16].

Store min-cost-path trees. For each leg of $\overline{\Delta}[\ell, v]$, we store pairs $\langle t_\ell, PRED[\ell, v](t_\ell)\rangle$ of departure-times t_ℓ from ℓ and the predecessor of v in the corresponding min-cost-path tree rooted at (ℓ, t_ℓ), omitting the actual min-travel-time values $D[\ell, v](t_\ell)$. This modification makes the oracle aware only of the min-cost-path-tree structures created during the repeated sampling procedure. Additionally, rather than storing repeatedly the IDs of predecessors, which would be space consuming in networks with millions of vertices, we only store the position of the corresponding arc in the list of incoming arcs to a vertex v. Since the maximum in-degree in the road instances we have at our disposal is at most 7, we only need to consume 1 byte per storage for a predecessor. We could even consume 3 bits per predecessor, which could then be packed into only two bytes containing also the corresponding departure-time value (by an appropriate discretization of the departure-time values). We prefer *not* to combine predecessors with departure-times in the same bit string, because we shall exploit later the extensive repetition of identical sequences of departure-times, which nevertheless would be lost for strings also containing the predecessors. It was observed in both benchmark instances that about one half of all possible destinations per landmark ℓ appear to have a *unique* predecessor throughout the entire period of departure-times, $[0, T)$. For them we store their unique predecessor only once. For the remaining destinations though, even with only two possible predecessors, we have to store the entire sequence of predecessor-changes.

Omit intermediate breakpoints. TRAP computes, and explicitly stores, intermediate breakpoints $(\overline{t}_m, \overline{D}_m)$ between consecutive sampled breakpoints of $D[\ell, v]$, as the intersection points of the two legs involved in the definition of $\overline{\delta}[\ell, v](t)$ (cf. Figure 1), for each pair (ℓ, v) and those intervals where the MAE is sufficiently small and v becomes deactivated. In CTRAP we choose *not* to keep these intermediate breakpoints and restrict the preprocessed information only to the actual samples. We let the query algorithm deal with the missing information, whenever needed. This way we avoid storing approximately 10M (for Berlin) and 100M (for Germany) of intermediate breakpoints per landmark.

Merge sequences of breakpoints with identical predecessors. CTRAP's next algorithmic intervention is based on the observation that the vast majority of all destinations appear to have on average 2 alternating predecessors throughout the entire period $[0, T)$. To save space, we choose to merge *consecutive* sampled breakpoints for v of the form $\langle t_\ell, x = PRED[\ell, v](t_\ell)\rangle$ and $\langle t'_\ell, x = PRED[\ell, v](t'_\ell)\rangle$, i.e., possessing the same predecessor. This leads to a reduction in the number of breakpoints to store, but also has a negative influence on the similarities of the departure-times sequences, and thus on the repetitions that we could avoid (see next heuristic). However, there is still positive gain by applying both this heuristic and that for avoiding multiple copies of departure-times sequences.

Avoid multiple copies of common departure-time sequences. CTRAP's next key insight is based on the fact that it is a repeated-sampling method which probes (at common departure-times for all destinations) min-cost-path trees from a landmark ℓ, starting from a coarse-grained sampling towards more fine-grained samples of the entire period $[0, T)$, until the MAE guarantee is satisfied for all reachable destinations from ℓ. A destination v may not care for all these departure-times, because the value of MAE may be satisfied at an early stage for it. This indeed depends on the actual minimum travel-time $\min\{D[\ell, v](t_s), D[\ell, v](t_f)\}$ at the endpoints of each given subinterval $[t_s, t_f)$. For each landmark-destination pair (ℓ, v),

we store the sequences $DEP[\ell, v]$ of necessary departure-times and $PRED[\ell, v]$ of the corresponding predecessors. The crucial observation is that destinations which are (roughly) at the same distance from ℓ are anticipated to have the same sequence of sampled departure-times, possibly differing only in their sequences of predecessors. It is clearly a waste of space to store two identical sequences $DEP[\ell, v] = DEP[\ell, u]$ more than once, even if the corresponding sequences of predecessors differ. Thus, we store each departure-times sequence as soon as it first appears for some destination v, and consider v as the *representative* of all other destinations u for which $DEP[\ell, u] = DEP[\ell, v]$. For each non-representative destination u, we store $PRED[\ell, u]$ and the corresponding representative v. Our next challenge is to efficiently compare departure-times sequences. To avoid a potential blow-up of the preprocessing time, we do not compare them point-by-point. Instead, we assign to every sampled departure-time t_ℓ two **iuar**[2] chosen floating-point numbers $w_1(t_\ell), w_2(t_\ell)$ from the interval $[1.0, 100.0]$. Each destination u adds the two values $w_1(t_\ell) \cdot t_\ell$ and $w_2(t_\ell) \cdot t_\ell$ to its own hash keys, i.e., $H_1[u] = H_1[u] + w_1(t_\ell) \cdot t_\ell$ and $H_2[u] = H_2[u] + w_2(t_\ell) \cdot t_\ell$, *only when* t_ℓ is indeed a necessary sample for u. Otherwise, the hash keys of u remain intact. At the end of the sampling process, we sort lexicographically the hash pairs of all destinations, in order to discover families of common departure-times sequences. We deduce that two destinations possess the same sequence when both their hash pairs match, in which case we verify this allegation by comparing them point by point. We observed that, for both benchmark instances, 80% of all destinations with at least two predecessors can be represented w.r.t departure-times by the remaining 20% of (representative) destinations.

Indexing preprocessed information. For retrieving efficiently the summaries from a landmark ℓ to each destination v, we maintain a vector of pointers per landmark, one pointer per destination, providing the address for the starting location of the summary for v. The pointers are in ascending order of vertex ID. The lookup time is $\mathcal{O}(1)$ and the required space for this indexing scheme is $\mathcal{O}(n \cdot |L|)$ additional bytes, where L is the chosen landmark set.

Speeding up preprocessing time. Handling only min-cost-path trees also has a collateral effect of speeding up the required preprocessing time. The reason for this is that we do not compute explicitly, each and every time that we sample travel-time values from ℓ, the exact shapes of the corresponding minimum-travel-time functions per destination. The travel-time summaries provided by `FLAT` were created based on this explicit computation of all the earliest-arrival *functions* per destination v, from each landmark ℓ. In contrast, the min-cost-path summaries of `CFLAT` are created without having to compute earliest-arrival functions. This leads to a reduction in the preprocessing time of more than 60%.

2.2.2 The Query Algorithm `CFCA(N)`

`CFCA(N)` is based on `FCA(N)` [15], but is fundamentally different from it in the sense that it exploits min-cost-path trees, and also considers the *od*-path construction as part of it, which was not the case for `FCA(N)`, and indeed for most of the query algorithms in the literature. N indicates the number of landmarks to be settled by `CFCA(N)` around the origin o. The pseudocode of the algorithm is presented in the next paragraph. `CFCA(N)` works as follows. In case that the destination d is already settled in Step 1, the resulting (exact) *od*-path can be computed by backtracking towards the origin, following the pointers to all predecessors. Otherwise, we proceed as follows. For each settled landmark ℓ, we have an optimal $o\ell$-path guaranteeing arrival-time $t_\ell = t_o + D[o, \ell](t_o)$ at ℓ. Since we do not have at our disposal travel-time values from ℓ towards d, or any other vertex, we are not able to compare ℓv-paths based on their (approximate) lengths. On the other hand, for the given departure-times t_ℓ

[2] **iuar** = independently and uniformly at random, without repetitions.

and any vertex v, we can tell the predecessor(s) of v in the (at most two per landmark) most relevant min-cost-path trees, the ones at the consecutive sampled departure-times t_ℓ^- and t_ℓ^+ of each $DEP[\ell, v]$ for which it holds that $t_\ell \in [t_\ell^-, t_\ell^+)$.

procedure `CFCA(N)`

STEP 1: A `TDD` ball is grown from (o, t_o), until N landmarks are settled.
1.1: **if** d is already settled **then return** optimal solution.
1.2: For each settled landmark ℓ, $t_\ell = t_o + D[o, \ell](t_o)$.

STEP 2: An appropriate subgraph is recursively created from d.
2.1: $Q = \{\ d\ \}$ /* Q is a FIFO queue */
2.2: **while** $\neg Q.Empty()$ **do** :
2.3: **if** $v = Q.Pop()$ is not explored from STEP 1's `TDD` ball **then** :
2.4: **for** each settled landmark ℓ of STEP 1 **do** :
2.5: Mark the arcs $\langle PRED[\ell, v](t_\ell^-), v\rangle$ and $\langle PRED[\ell, v](t_\ell^+), v\rangle$ leading to v, where $[t_\ell^-, t_\ell^+)$ is the unique interval in $DEP[\ell, v]$ containing t_ℓ.
2.6: $Q.Push(PRED[\ell, v](t_\ell^-))$; $Q.Push(PRED[\ell, v](t_\ell^+))$
2.7: **end for**
2.8: **end while**

STEP 3: **return** optimal od-path in the induced subgraph by (`TDD` ball of) STEP 1 and STEP 2.

`CFCA(N)` marks (per settled landmark ℓ) the connecting arcs from these most relevant predecessor(s) $PRED[\ell, v](_\ell^-)$ and $PRED[\ell, v](_\ell^+)$, towards v. All these discovered predecessors w.r.t. the N settled landmarks are inserted (if not already there) in a FIFO queue, which was initialized with d, so that, upon their extraction from the queue, they can provide in turn their own predecessors, etc. The recursive search for predecessors stops as soon as a vertex x in the explored area of the initial `TDD` ball of Step 1 is reached. `CFCA` marks then also the arcs of the corresponding short (not necessarily the shortest though, since x is explored but not necessarily settled) ox-path. This way we are guaranteed that in the subgraph of marked arcs there is already an od-path which has been oriented by (ℓ, t_ℓ) and passes via x. Step 2 of `CFCA(N)` terminates when the FIFO queue becomes empty, i.e., we no longer have to process intermediate vertices which are unexplored by Step 1. The actual path construction takes place in Step 3, which considers the subgraph induced by the marked arcs and continues growing the `TDD` ball from (o, t_o) within this subgraph. This path construction indeed leads to significantly smaller relative errors, since the resulting od-path is not only the best *prediction* among a given set of N paths induced by the N settled landmarks (as in `FLAT`), but actually the *optimal* od-path within the induced sugbgraph.

The worst-case approximation guarantee of `CFCA`(1) is $(1+\varepsilon+\psi)$ (identical to that of `FCA` [16]), where ε is `CTRAP`'s approximation guarantee and ψ is a constant depending on ε and the travel-time metric (but not on the size) of the network. Note that we could theoretically improve the stretch of `CFCA(N)` to $(1+\sigma)$, for any constant $\sigma > \varepsilon$, and get a PTAS, by using in Step 1 the `RQA` algorithm [16]. We choose *not* to do so, because our previous experimental evaluation with `FLAT` [15] has shown that `FCA(N)` in practice dominates `RQA`.

3 Experimental Evaluation

Experimental Setup and Goal. Our algorithms were implemented in C++ (GNU GCC version 5.4.0) and Ubuntu Linux (16.04 LTS). All the experiments were conducted on a 6-core Intel(R) Xeon(R) CPU E5-2643v3 3.40GHz machine, with 128GB of RAM. We used 12 threads for the parallelization of the preprocessing phase. `CFCA` was always executed on a single thread. For the sake of comparison, we used the same set of 50,000 queries, **iuar** chosen from $V \times V \times [0, T)$ in each instance, for all possible landmark sets. The PGL library [18] was used for graph representation and operations. Two benchmark instances were used, the first concerning the city of Berlin, and the second the national road network of Germany.

More details on the availability of code and data are provided in Appendix C.

The main goal of our experimental evaluation was to investigate the scalability of `CFLAT`: how smoothly does it trade higher preprocessing requirements for better approximation guarantees and query-times. To demonstrate this, we aim at showcasing the performance of `CFCA(N)` for several types and sizes of landmark sets. We also choose to increase the typical size of the used landmark sets in our comparison of different landmark selection policies.

Landmark Selection Policies. Although the preprocessing requirements are proportional to $|L|$ (number of landmarks), they are essentially invariant of the landmark selection policy. However, as previous experimental evaluation indicated [15], the performance of the query algorithms has a strong dependence on the type of the landmarks. A key observation was that the *sparsity* of landmarks (not being too close to each other) as well as their *importance*, are crucial parameters. Therefore, in this work we insist in almost all cases (except for the RANDOM landmark sets which are used as baseline) on selecting the landmarks sparsely throughout the network. As for their importance, when such information is available, we also consider the selection of landmarks at junctions of an important road segment (as in [15]). Finally, we consider a new measure of vertex significance, the (approximate) betweeness-centrality measure. In particular, we consider the following landmark selection policies:

◇ RANDOM (R): **iuar** choice of landmarks.

◇ SPARSE-RANDOM (SR): Incremental **iuar** choice of landmarks, where each chosen landmark excludes a free-flow neighborhood of vertices around it from future landmark selections.

◇ IMPORTANT-RANDOM (IR): A variant of R which moves each random landmark to its nearest important vertex within a free-flow neighborhood of size 100. This policy is only applicable for the instance of Berlin which provides road-segment importance information.

◇ SPARSE-KAHIP (SK): We use the `KaFFPa` algorithm of the KAHIP partitioning software (v1.00) [1], setting the parameters so that there are many more boundary vertices than the required number of landmarks. The landmarks are incrementally and **iuar** chosen among the boundary vertices. Each landmark excludes a free-flow neighborhood from future selections.

◇ KAHIP-CELLS (KC). Starting with a KAHIP partition, one landmark per cell is incrementally and **iuar** chosen, excluding a free-flow neighborhood from future selections.

◇ BETWEENESS-CENTRALITY (BC): Vertices are ordered in non-increasing *approximate betweeness-centrality* (ABC) values [2]. Landmarks are selected incrementally according to ABC values, excluding a free-flow neighborhood from future selections.

◇ KAHIP-BETWEENESS (KB): For a KAHIP partition, incrementally choose as landmark the vertex with the highest ABC value in a cell, excluding a neighborhood from future selections.

We finally consider the following systematic naming of the landmark sets. Each set is encoded as XY, where $X \in \{R, SR, IR, SK, KC, BC, KB\}$ determines the type of landmark set, and $Y \in \{250, 500, 1K, 2K, 3K, 4K, 8K, 16K, 32K\}$ determines its size.

Evaluation of `CFLAT` @ Berlin. For Berlin we have considered all types of landmarks. For each of them, we have used as baseline the size $Y = 4K$. $\{R, SR, IR, SK\}$ were considered also in [15] (but for smaller sizes), whereas $\{KC, BC, KB\}$ are new types. Especially for R we tried all possible values for Y, in order to showcase the scalability of `CFLAT` and its smooth trade-off of preprocessing requirements, query-times and stretch factors. Concerning *vertex-importance* (only available in Berlin), we considered as important those vertices which are incident to roads of category at most 3. As for *sparsity*, we set the sizes of the excluded free-flow ball per selected landmark to 150 vertices for SR, 100 for IR, 50 for SK, 20 for KC, 150 for BC, and 20 for KB. For KAHIP based landmark sets (SK, KC and KB) we used the following parameters: The number of cells to partition the graph was set to $4,000$, having $13,256$ boundary vertices in total. For SK we chose randomly $4,000$ boundary vertices as landmarks. For KC and KB we chose one landmark per cell.

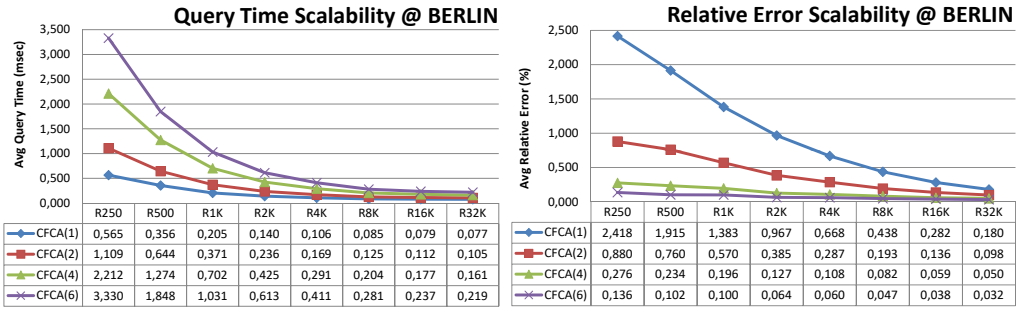

	R250	R500	R1K	R2K	R4K	R8K	R16K	R32K
CFCA(1)	0,565	0,356	0,205	0,140	0,106	0,085	0,079	0,077
CFCA(2)	1,109	0,644	0,371	0,236	0,169	0,125	0,112	0,105
CFCA(4)	2,212	1,274	0,702	0,425	0,291	0,204	0,177	0,161
CFCA(6)	3,330	1,848	1,031	0,613	0,411	0,281	0,237	0,219

	R250	R500	R1K	R2K	R4K	R8K	R16K	R32K
CFCA(1)	2,418	1,915	1,383	0,967	0,668	0,438	0,282	0,180
CFCA(2)	0,880	0,760	0,570	0,385	0,287	0,193	0,136	0,098
CFCA(4)	0,276	0,234	0,196	0,127	0,108	0,082	0,059	0,050
CFCA(6)	0,136	0,102	0,100	0,064	0,060	0,047	0,038	0,032

Figure 2 Performance of CFCA(N) in Berlin, for random landmarks and 50,000 *random queries*.

We first conducted an experiment to test the scalability of CFCA's performance as a function of N and the number of landmarks, always for R-type landmarks. As is evident from Figure 2, the average errors decrease linearly and the query-times decrease quadratically, as we double the number of landmarks. Additionally, notable "quick-and-dirty" answers are possible with only 250 landmarks, which require total space 0.7GiB (0.17GiB after compression), cf. Figure 6. In particular, the query performance (average query time and relative error) varies from 0.565msec and 2.418% ($N = 1$), to 3.330msec and 0.136% ($N = 6$). If query time is the main goal, then for BC8K+R8K, the query performance of CFCA varies from 0.076msec and 0.19% ($N = 1$), to 0.226msec and 0.022% ($N = 6$). Since the average query-time for TDD is 107.466msec[3], the achieved speedup is more than 1,414.

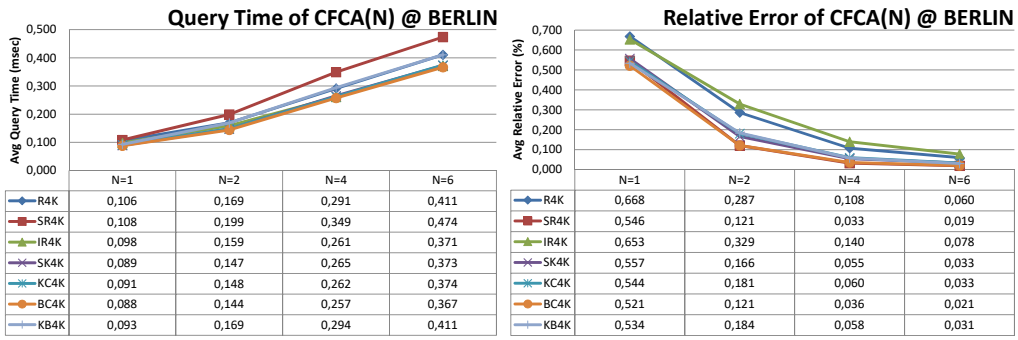

	N=1	N=2	N=4	N=6
R4K	0,106	0,169	0,291	0,411
SR4K	0,108	0,199	0,349	0,474
IR4K	0,098	0,159	0,261	0,371
SK4K	0,089	0,147	0,265	0,373
KC4K	0,091	0,148	0,262	0,374
BC4K	0,088	0,144	0,257	0,367
KB4K	0,093	0,169	0,294	0,411

	N=1	N=2	N=4	N=6
R4K	0,668	0,287	0,108	0,060
SR4K	0,546	0,121	0,033	0,019
IR4K	0,653	0,329	0,140	0,078
SK4K	0,557	0,166	0,055	0,033
KC4K	0,544	0,181	0,060	0,033
BC4K	0,521	0,121	0,036	0,021
KB4K	0,534	0,184	0,058	0,031

Figure 3 Performance of CFCA(N) in Berlin, for $4K$ landmarks and 50,000 *random queries*.

Our next experiment compares landmark types of size 4K each (cf. Figure 3). Concerning query-times, the best curve is that of BC4K. As for relative errors, SR4K and BC4K are clear winners. Further experiments are reported in Section D. In comparison with FLAT, the query-performance of CFCA(1) for BC4K is comparable (0.088msec and 0.521%) to that of FCA(1) (0.081msec and 0.771%) in [15]. We also tested hybrid landmark sets. Interestingly, we achieved our best query performance with the hybrid set BC8K+R8K, which varies from 0.076msec and 0.192% (for $N = 1$), to 0.226msec and 0.022% (for $N = 6$). It is also observed that, as we mix BC-landmarks with R-landmarks, the more BC landmarks we get the better for the relative error, whereas query-time is favored by more R-landmarks (cf. Figure 7).

[3] TDD is executed here on the original instance, even before the vertex contraction. In [15] it was executed on the contracted graph, hence the slightly smaller execution times of TDD in that work. Nevertheless, we believe that this is the appropriate measurement to make for TDD, for sake of comparison with other works, and also since the contraction of degree-2 vertices is part of the preprocessing phase.

Evaluation of CFLAT @ Germany. We considered R-landmark sets of sizes from 1K to 4K. The rest of the landmark sets were of size 3K, with excluded neighborhood size 1,200 vertices for SR3K, 350 for SK3K, and 1,000 for BC3K. We started again with a demonstration of the scalability of CFCA on R-landmark sets, as a function of the number of landmarks (cf. Figure 4). The relative errors decrease linearly and the running times decrease quadratically, as we increase the number of landmarks. Remarkable relative errors of 0.071% are achieved for CFCA(6) even with 1K landmarks which require 26.8GiB (8.1GiB compressed) space, with query-time 11.974msec. Moreover, a "quick-and-dirty" answer of error at most 1.582% is returned in only 2.175msec. The best query-times and relative errors are achieved for R4K, where CFCA(1) achieves 0.819msec and 0.911%, and CFCA(6) has 4.201msec and 0.049%.

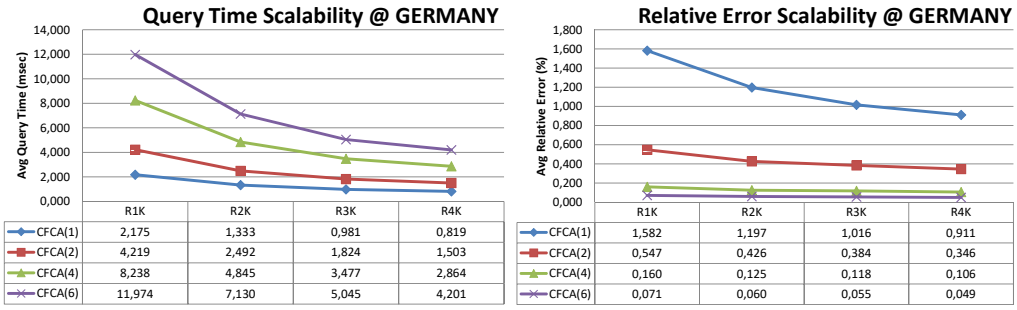

	R1K	R2K	R3K	R4K
CFCA(1)	2,175	1,333	0,981	0,819
CFCA(2)	4,219	2,492	1,824	1,503
CFCA(4)	8,238	4,845	3,477	2,864
CFCA(6)	11,974	7,130	5,045	4,201

	R1K	R2K	R3K	R4K
CFCA(1)	1,582	1,197	1,016	0,911
CFCA(2)	0,547	0,426	0,384	0,346
CFCA(4)	0,160	0,125	0,118	0,106
CFCA(6)	0,071	0,060	0,055	0,049

Figure 4 Performance of CFCA(N) in Germany, for random landmarks and 50,000 *random queries*.

We proceeded next with a comparison of various landmark types of size 3,000 each (cf. Figure 5). For Germany we have a clear winner, BC3K, w.r.t. both query-times and relative errors and $N \leq 2$. For $N \in \{4, 6\}$, SK3K is the fastest and SR3K is the most accurate landmark policy. Since the average time of TDD is 1,421.12msec, the best speedup for 3K landmarks is 1,938, and the corresponding error is 0.911%. Once more, the best query performance is achieved by a hybrid landmark set. In particular, for BC3K+R1K CFCA's performance varies from 0.683msec and 0.831% (for $N = 1$), to 4.104msec and 0.031% (for $N = 6$), see Figure 8. Further experiments are reported in Section D.

Comparison with State-Of-Art. Table 1 presents a comparison with the most competitive speedup heuristics and oracles for TDRP. Details are provided in Section F. We compare the performances of the following algorithms, on the instances of Berlin and Germany: **(1)** TDCRP, tested on a 16-core Intel Xeon E5-2670 clocked at 2.6 GHz, with 64GB of DDR3-1600 RAM, 20 MB of L3 and 256 KB of L2 cache. The reported numbers are from [4]; **(2)** FreeFlow, TD-S and TD-S+A, tested on a 16-core Intel Xeon E5-1630 v3 clocked at 3.70GHz

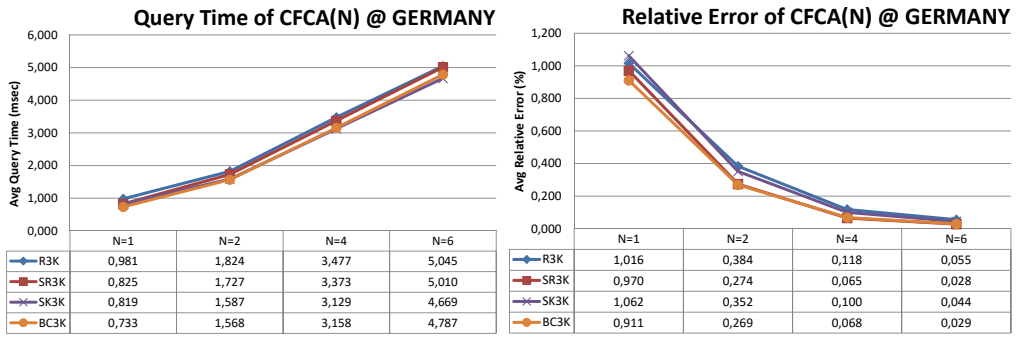

	N=1	N=2	N=4	N=6
R3K	0,981	1,824	3,477	5,045
SR3K	0,825	1,727	3,373	5,010
SK3K	0,819	1,587	3,129	4,669
BC3K	0,733	1,568	3,158	4,787

	N=1	N=2	N=4	N=6
R3K	1,016	0,384	0,118	0,055
SR3K	0,970	0,274	0,065	0,028
SK3K	1,062	0,352	0,100	0,044
BC3K	0,911	0,269	0,068	0,029

Figure 5 Performance of CFCA(N) in Germany, for 3K landmarks and 50,000 *random queries*.

with 128GB of 2133GHz DDR4 RAM. The reported numbers are from [22]; **(3)** `inex.TCH`,

	Algorithm		Preprocessing Performance			Query Performance			
	Name [ref.]	Parameters	Time h:m (#cores)	Work h:m	Space B/node	Path N/Y	Time msec	error (%) avg	max
GERMANY	TDD [⋆]	–	–	–	–	•	1,421	0	0
	inex.TCH [3]	(0.1)	06:18 (8)	50:24	286	○	0.70	0.02	0.10
		(1.0)			214		0.69	0.27	1.01
		(2.5)			172		0.72	0.79	2.44
		(10.0)			113		1.06	3.84	9.75
	KaTCH [⋆]	–	34:29 (6)	206:56	9.029	○	26.9	25.6	1245.06
	TDCRP [4]	(1.0)	00:13 (16)	03:28	77	○	1.17	0.68	3.60
	FreeFlow [22]	–	00:07 (16)	01:57	n/r	•	0.24	0.031	2.516
	TD-S [22]						0.6	0.000746	0.989
	TD-S+A [22]						6.36	0.000312	0.227
	DijFreeFlow [⋆]	–	–	–	–	•	736.24	0.352	17.569
	FLAT [15]	SR2K, N=1	42:42 (6)	256:12	106,075	○	1.275	1.444	n/r
		SR2K, N=6					9.952	0.662	
		SK2K, N=1	44:06 (6)	264:36			1.269	1.534	
		SK2K, N=6					9.689	0.676	
	CFLAT [⋆]	BC4K, N=1	32:36 (6)	195:29	30,769	•	0.693	0.858	19.154
		BC4K, N=6					3.841	0.049	
		BC3K+R1K, N=1	32:36 (6)	195:29	30,769		**0.683**	0.831	
		BC3K+R1K, N=6					4.104	**0.031**	
BERLIN	TDD [⋆]	–	–	–	–	•	107.5	0	0
	KaTCH [⋆]	–	< 00:01 (6)	< 00:04	0.593	○	0.3	0.41	47.74
	TDCRP [4]	(1.0)	00:02 (16)	00:28	67	○	0.28	1.47	2.69
	FreeFlow [22]	–	< 00:01 (16)	00:07	n/r	•	0.09	0.0165	1.343
	TD-S [22]						0.23	0.00022	0.254
	TD-S+A [22]						3.01	0.000086	0.158
	DijFreeFlow [⋆]	–	–	–	–	•	54.608	0.367	20.42
	FLAT [15]	SR2K, N=1	05:12 (6)	31:12	48,389	○	0.081	0.771	n/r
		SR2K, N=6					0.586	0.317	
		SK2K, N=1	05:42 (6)	33:12	52,826		0.083	0.781	
		SK2K, N=6					0.616	0.227	
	CFLAT [⋆]	BC4K, N=1	03:44 (6)	22:23	6,353	•	0.088	0.521	16.167
		BC4K, N=6					0.367	0.021	
		BC16K, N=1	14:42 (6)	88:12	27,226		0.078	0.227	10.063
		BC16K, N=6					0.250	0.019	
		BC8K+R8K, N=1					**0.076**	0.192	
		BC8K+R8K, N=6					0.226	**0.022**	

Table 1 Comparison with State-Of-The-Art.

tested on an 8-Core Intel i7, clocked at 2.67 GHz, with 64 GB DDR4 RAM. The reported numbers are from [4]; **(4)** an open-source version of `TCH` (`KaTCH`[4]), tested (with compilation parameters -O3 and -DNDEBUG, and its default values) on our machine; **(5)** our own implementation of the `FreeFlow` heuristic (called `DijFreeFlow`), tested on our machine (it is a static-Dijkstra execution on the Free Flow instance, with no exploitation of any speedup heuristic, and then computation of the time-dependent travel-time along the chosen path); and **(6)** `FLAT` and `CFLAT`, which were tested on our machine. The reported numbers for `FLAT` are from [15]. All the reported times are *unscaled* (i.e., as they have been reported) and include both metric-independent and metric-dependent preprocessing of the instances. *Work* is measured as the product of the running time with the number of cores. The "path" column indicates whether the explicit construction of a connecting path is accounted for in the reported query times. ○ is a NO-answer, • means YES. "n/r" means that a particular value has not been reported. The algorithms `TDD`, `KaTCH`, `DijFreeFlow` and `CFLAT`, marked in Table 1 with [⋆], were evaluated in the present work, on exactly the same benchmark instances and for the same sets of 50K **iuar** chosen queries.

[4] `https://github.com/GVeitBatz/KaTCH`, with checksum 70b18ad0791a687c554fbfe9039edf79bc3a8ff3.

Acknowledgements

The authors wish to thank G. Veit Batz, Julian Dibbelt and Ben Strasser for valuable and fruitful discussions.

References

1 KaHIP – Karlsruhe High Quality Partitioning, May 2014.
2 D. Bader, S. Kintali, K. Madduri, M. Mihail: Approximating betweenness centrality. *Algorithms and Models for the Web-Graph (WAW)*, pp. 124–137, Springer (2007)
3 G. V. Batz, R. Geisberger, P. Sanders, C. Vetter: Minimum time-dependent travel times with contraction hierarchies. *J. of Experimental Algorithmics*, **18**(1.4):1–43, ACM (2013).
4 M. Baum, J. Dibbelt, T. Pajor, D. Wagner: Dynamic time-dependent route planning in road networks with user preferences. *Experimental Algorithms (SEA)*, **LNCS 9685**:33–49, Springer (2016)
5 F. Dehne, M. T. Omran, and J.-R. Sack. Shortest paths in time-dependent FIFO networks. *Algorithmica*, **62**(1-2):416–435 (2012)
6 D. Delling: Time-Dependent SHARC-Routing. *Algorithmica*, **60**(1):60–94 (2011)
7 D. Delling, G. Nannicini: Core routing on dynamic time-dependent road networks. *Informs J. on Computing*, **24**(2):187–201 (2012)
8 D. Delling, D. Wagner: Landmark-based routing in dynamic graphs. *Experimental Algorithms (WEA'07)*, **LNCS 4525**:52–65, Springer (2007)
9 D. Delling, D. Wagner: Time-dependent route planning. *Robust and Online Large-Scale Optimization*, **LNCS 5868**:207–230, Springer (2009)
10 U. Demiryurek, F. Banaei-Kashani, C. Shahabi: A case for time-dependent shortest path computation in spatial networks. *SIGSPATIAL Advances in Geographic Information Systems (GIS)*, pp. 474–477, ACM (2010).
11 S. E. Dreyfus: An appraisal of some shortest-path algorithms. *Operations Research*, **17**(3):395–412, 1969.
12 L. Foschini, J. Hershberger, S. Suri: On the complexity of time-dependent shortest paths. *Algorithmica*, **68**(4):1075–1097 (2014).
13 M. Hilger, E. Köhler, R. H. Möhring, H. Schilling: Fast point-to-point shortest path computations with arc-flags. *The Shortest Path Problem: Ninth DIMACS Implementation Challenge*, **74**: 41–72, AMS (2009)
14 S. Kontogiannis, G. Michalopoulos, G. Papastavrou, A. Paraskevopoulos, D. Wagner, C. Zaroliagis: Analysis and experimental evaluation of time-dependent distance oracles. *Algorithm Engineering and Experiments (ALENEX)*, pp. 147–158, SIAM (2015)
15 S. Kontogiannis, G. Michalopoulos, G. Papastavrou, A. Paraskevopoulos, D. Wagner, C. Zaroliagis: Engineering oracles for time-dependent road networks. *Algorithm Engineering and Experiments (ALENEX)*, pp. 1–14, SIAM (2016)
16 S. Kontogiannis, D. Wagner, C. Zaroliagis: Hierarchical oracles for time-dependent networks. *Algorithms and Computation (ISAAC)*, **64**(47):1-13, LIPICS (2016)
17 S. Kontogiannis, C. Zaroliagis: Distance oracles for time-dependent networks. *Algorithmica*, **74**(4):1404–1434 (2016)
18 G. Mali, P. Michail, A. Paraskevopoulos, C. Zaroliagis: A new dynamic graph structure for large-scale transportation networks. *Algorithms and Complexity (CIAC)*, **LNCS 7878**:312–323, Springer (2013)
19 G. Nannicini, D. Delling, L. Liberti, D. Schultes: Bidirectional A* search on time-dependent road networks. *Networks*, **59**:240–251 (2012)
20 M. Omran and J.-R. Sack: Improved approximation for time-dependent shortest paths. *Computing and Combinatorics (COCOON)*, **LNCS 8591**:453–464, Springer (2014)
21 A. Orda, R. Rom: Shortest-path and minimum delay algorithms in networks with time-dependent edge-length. *J. of the ACM*, **37**(3):607–625, ACM (1990)
22 B. Strasser: Intriguingly simple and efficient time-dependent routing in road networks. ArXiv technical report (arxiv:1606.06636v1), Karlsruhe Institute of Technology (2016)

A Preprocessing the Instances

We recap at this point some heuristic improvements which are inherited from FLAT towards simplifying the road instance and thus saving space.

Contraction of the road network. The preprocessing space and time can be reduced if we only focus on a subgraph of the underlying graph representing the road network. Towards this direction, we have chosen to "contract" all the vertices which do not depict junctions of road segments (e.g., intermediate stops along a road segment). We consider these vertices as *inactive* (only for the preprocessing phase), and we do not consider them during the subsequent preprocessing of travel-time related information, since they do not provide actual alternatives along a route using them, unless they are indeed endpoints of the query at hand. It is emphasized though, that the queries are conducted in the original graph, not just the contracted subgraph, meaning that we can query also for contracted origin-destination pairs and the returned paths do not contain shortcuts but actual road segments.

In more detail, in the *instance-contraction phase* we seek for maximal w.r.t. the number of arcs (possibly bidirectional) paths which have no "vertical" intersections, i.e., all the intermediate vertices connect only with their neighboring vertices along the path. Each such path is substituted with a *shortcut* (arc) connecting its endpoints, which is equipped with an arc-travel-time function equal to the corresponding exact path-travel-time function. In fact, multiple paths with no intermediate intersections may connect the same active endpoints. In that case, a single shortcut represents more than one contracted paths, i.e. the arc-travel-time function of the shortcut is computed by applying the minimization operator on the path-travel-time functions corresponding to each of the contracted paths. If there exists an original arc connecting two active endpoints, which are to be connected with a shortcut, we choose not to insert an additional shortcut, but to update accordingly the arc-travel-time of the already existing arc which now plays the role of a shortcut as well. The original arcs involved in the contracted paths are also considered as *inactive*. All contracted vertices are ignored during the landmark-preprocessing and therefore the number of reachable destinations from a landmark is smaller. At the query phase, the contracted paths can be easily recovered, by exploiting the appropriate information kept on all shortcuts and the corresponding contracted vertices.

Almost constant legs. The original TRAP approximation method [15] introduced at least one intermediate breakpoint per interval that does not yet meet the required approximation guarantee. This is certainly unnecessary for small intervals in which the actual shortest-travel-time functions are constant. To avoid the blow-up of the required preprocessing space, we heuristically make a "guess" that we have to deal with a constant shortest-travel-time function $D[\ell, v]$ within a given interval $[t_s, t_f = t_s + \tau)$ with sufficiently small length τ, whenever the following holds: $D[\ell, v](t_s) = D[\ell, v](t_f) = D[\ell, v]\left(\frac{t_s+t_f}{2}\right)$. This is justified by the fact that $D[\ell, v]$ is a continuous pwl function and it is unlikely that three different departure-times within a small interval would give the same value, unless the function is indeed constant. Of course, one could easily construct artificial examples for which this criterion is violated, e.g. by providing a properly chosen periodic function with period $\tau/2$. On the other hand, one can easily tackle this by considering a *randomly perturbed* sampling period $\tau + \delta$, for some arbitrarily small but positive random variable δ. Since we engineer oracles for real-world road-networks, having three colinear points which do not belong to a leg of the sampled travel-time function is quite unlikely, therefore we choose not to randomly

perturb the sampling period.

Fixed range. For a one-day time period, departure-times and arrival-times have a bounded value range. The same also holds for travel times which are at most one-day for any query within a country area such as Germany. Therefore, when the considered precision of the traffic data is within seconds, we handle time-values as integers in the range $\{0\ ,\ 1\ ,\ \ldots\ ,\ 86,399\}$, for milliseconds as integers in $\{0\ ,\ 1\ ,\ \ldots\ ,\ 86,399,999\}$, etc.

Any (real) time value within a single-day period, represented as a floating-point number t_f, can thus be converted to an integer t_i with fewer bytes and a given unit of measure. For a unit measure (or scale factor) s, the resulting integer is $t_i = \left\lceil \frac{t_f}{s} \right\rceil$, requiring $\left\lceil \frac{\log_2(t_f/s)}{8} \right\rceil$ bytes for its storage. The division $\frac{t_f}{s}$ has quotient π and remainder υ s.t., $t_f = s \cdot \pi + \upsilon$, and $t_i = \left\lceil \frac{s\cdot\pi+\upsilon}{s} \right\rceil = \pi + \left\lceil \frac{\upsilon}{s} \right\rceil \in \{\pi, \pi+1\}$, since $0 \leq \upsilon \leq s-1$. Therefore, by storing t_i we actually consider the upper-approximating time $t'_f = s \cdot t_i$ of t_f, which causes an absolute error of at most s (i.e., one unit of measure): $t'_f - t_f < s \cdot (\pi+1) - s \cdot \pi = s$. In our experiments, for storing the time values involved in the approximate shortest-travel-time functions, we have considered a $1.32sec$ resolution, corresponding to the appropriate scale factor $s = 1.318359375$ (when originally counting time in seconds), that requires 2 bytes per time-value.

B The CTRAP approximation algorithm (pseudocode)

We now present a more detailed description of CTRAP. We start with the data types used in by the algorithm. For a given landmark vertex ℓ, a destination vertex v, and a subinterval $[t_s, t_f) \subseteq [0, T)$, the flag $ACTIVE[\ell, v](t_s, t_f)$ declares whether the upper-approximation $\overline{\delta}[\ell, v]$ considered by CTRAP (cf. Figure 1) is satisfactory, given the required approximation guarantee that we consider. The variable τ determines the current step of the sampled departure-times from ℓ. $PRED[\ell, v]$ and $DEP[\ell, v]$ are the sequences of predecessors and (corresponding) departure-times from ℓ, w.r.t. the destination vertex v. We assure that $DEP[\ell, v]$ is always ordered in increasing departure-time values. This is done by assuming the operation $DEP[\ell, v].SortedInsertion(x)$ which places x in the right position, which is then returned by the procedure. As for $DEP[\ell, v]$, we consider the insertion of a new element u at an arbitrary position pos, $DEP[\ell, v](u, pos)$. It is mentioned at this point that these operations have been implemented in a rather straightforward manner (essentially performing linear scans on the queues), leaving for the future the consideration of more sophisticated implementations.

The boolean function $MAE[\ell, v](t_s, t_f)$ determines whether the maximum-absolute-error test is satisfied for v, in the interval $[t_s, t_f)$. In particular, since we already have sampled all the travel-times at t_s and t_f, for a given approximation guarantee $\varepsilon > 0$ we perform the following test, which is a sufficient condition for $\overline{\delta}[\ell, v]$ being a $(1+\varepsilon)$-upper-approximation of $D[\ell, v]$ within $[t_s, t_f)$:

procedure MAE$[\ell, v](t_s, t_f, \varepsilon)$
1: **if** $\min\{D[\ell, v](t_s), D[\ell, v](t_f)\} \geq \left(1 + \frac{1}{\varepsilon}\right) \Lambda_{\max}$ **then** **return** (TRUE)
2: **else** **return** (FALSE)

The pseudocode of CTRAP is the following:

procedure CTRAP(ℓ, ε)

for $v \in V$ **do** { $ACTIVE[\ell, v](0, T) = TRUE$ }; $\tau_{old} = T$; $\tau = 3200$ /* initialization */
while $\exists v \in V, \exists k \in [0, T) : \ ACTIVE[\ell, v](k\tau_{old}, (k+1)\tau_{old}) == TRUE$ **do**
  Sample min-cost-path trees rooted at ℓ, only for *new* departure-times $k\tau \in [0, T)$
  /* $(w_1(k\tau), w_2(k\tau))$ is the pair of random seeds for $t_\ell = k\tau$. */
  /* $PRED[\ell, v](k\tau)$ indicates v's parent in the tree routed at $(\ell, k\tau)$. */
  for $v \in V \wedge k : k\tau \in [0, T)$ **do** /* looking for still active destinations... */
    if $ACTIVE[\ell, v](k\tau_{old}, (k+1)\tau_{old}) == TRUE$ **then**
      $HASH[v] = HASH[v] + (w_1(k\tau), w_2(k\tau)) \cdot k\tau$ /* Update hash keys... */
      if $DEP[\ell, v].NotInSequence(kt)$ **then**
        $position = DEP[\ell, v].SortedInsertion(k\tau)$;
        $PRED[\ell, v].Insertion(parent[\ell, v](k\tau), position)$
      end if
      if $MAE[\ell, v](k\tau, (k+1)\tau, \varepsilon) == TRUE$ **then** { $ACTIVE[\ell, v](k\tau, (k+1)\tau) = FALSE$ }
    end if
  end for
  $\tau = \tau/2$; $\tau_{old} = 2\tau$
end while
for $v \in V$ **do**
  repeat /* merge intervals with the same predecessor... */
    for *consecutive* records $(PRED[\ell, v](t), t)$ and $(PRED[\ell, v](t'), t')$ such that $PRED[\ell, v](t) == PRED[\ell, v](t')$ **do**
      $PRED[\ell, v].Delete(PRED[\ell, v](t'))$
      $DEP[\ell, v].Delete(t')$
    end for
  until $PRED[\ell, v]$ does not have identical consecutive records.
end for
Lexicographically sort in $DEST[\ell]$ the destinations v according to their hash key pairs.
for $v \in DEST[\ell]$ (in the previous lex-order) **do** /* Avoid multiple copies of dep-time sequences... */
  if $HASH[v] == HASH[DEST[\ell].Previous[v]]$
  then { $representative[v] = DEST.Previous[v]$; $DEP[\ell, v].Destroy()$ }
end for

C Benchmark Instances and Preprocessing

Our implementations and data sets constitute part of a broader route planning application service developed within the frame of EU-funded projects, which has been piloted in the cities of Berlin, Vitoria and Athens, as well as in the national road network of Germany. Due to complicated IPR issues, we cannot make our source code and benchmark data publicly available.

We proceed in this section with a detailed presentation of the benchmark instances of Berlin and Germany, on which we have conducted the experimental evaluation of CFLAT.

Berlin Instance. The instance of Berlin (kindly provided by TomTom in the frame of common R&D projects) consists of $473,253$ nodes and $1,126,468$ arcs. The instance-preprocessing heuristic A created $183,468$ shortcuts. Whenever more than one contracted paths shared the same endpoints, we added only one shortcut representing all these contracted paths. There were 914 such cases in the Berlin instance. The contracted paths that could be represented by an original arc in the graph, are $11,398$ in total. In overall, the contraction of Berlin led to a graph of $292,356$ active vertices and $752,362$ active arcs.

Germany Instance. The instance of Germany (kindly provided by PTV AG in the frame of common R&D projects) consists of $4,692,091$ nodes and $11,183,060$ arcs. After the instance-preprocessing phase we got an instance with $3,431,213$ active vertices and $11,554,840$ active arcs. The total number of the added shortcuts was $4,595,148$. We avoided the insertion of additional shortcuts in $106,464$ cases, where $6,816$ of them correspond to "parallel"

shortcuts and the 99,648 correspond to the existence of actual arcs connecting the endpoints of contracted paths.

Statistics for Berlin and Germany Instances. Table 2 reports some significant preprocessing statistics for the two instances. In particular, the measurements are the following: *(i)* the average number of vertices per landmark whose predecessor remains constant on the min-cost-path tree throughout the whole time period, *(ii)* the remaining vertices with pwl behaviour w.r.t. their predecessor, *(iii)* the average number of unique departure-time sequences stored, instead keeping one sequence per destination with pwl predecessor, and *(iv)* the average number of intermediate points of `TRAP` per landmark, which we now avoid to store.

	Vertices with Unique Pred	Vertices with pwl Pred	Unique Departure Time Sequences	Intermediate Points of `TRAP`
R4K	272,286	20,070	5,963	10,663,125
SR4K	272,287	20,069	5,831	10,688,275
IR4K	272,284	20,072	5,781	10,672,869
SK4K	272,282	20,074	6,011	10,934,712
KC4K	272,287	20,069	5,857	10,758,955
BC4K	272,293	20,063	5,858	10,728,776
KB4K	272,300	20,056	5,432	10,643,285

Table 2 Preprocessing statistics for `CFLAT` Oracle for Berlin.

Table 3 provides the preprocessing statistics related to Germany, in the same format as in the case of Berlin.

	Vertices with Unique Pred	Vertices with pwl Pred	Unique Departure Time Sequences	Intermediate Points of `TRAP`
R3K	3,201,577	229,636	38,102	112,137,488
SR3K	3,201,642	229,571	37,212	112,081,032
SK3K	3,201,503	229,710	38,068	113,536,811
BC3K	3,201,637	229,576	37,207	112,067,442

Table 3 Preprocessing statistics for `CFLAT` in Germany.

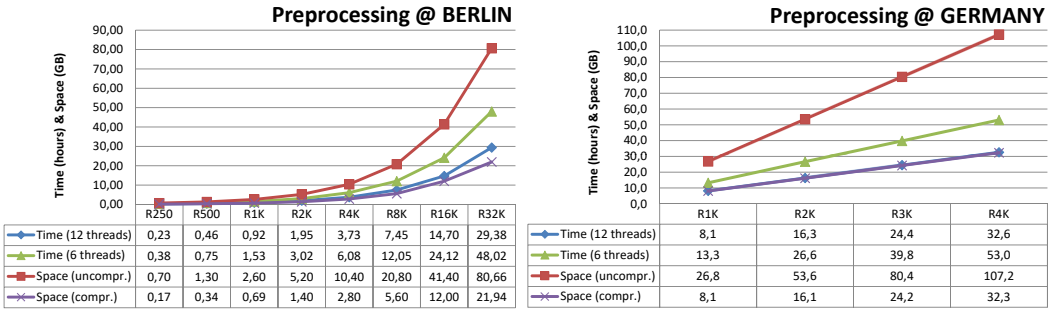

Figure 6 Preprocessing requirements for Berlin and Germany.

Preprocessing Requirements @ Berlin. We present in this section the preprocessing requirements for the construction of the summaries for `CFLAT`, for various sizes of random (R)

landmark sets (cf. Figure 6). The requirements for other landmark types are analogous. For this preprocessing, we have used 12 parallel threads on our 6-core machine.

It is worth mentioning that FLAT [15] required *uncompressed* preprocessing space 43GB, or equivalently, *compressed* size of 14MB per landmark, and 33h to preprocess R2K. On the contrary, with CFLAT R32K is preprocessed in 29.38h consuming 80.67GB (21.94GB compressed) space. As for R2K, it is preprocessed in 117min consuming only 5.2GB (1.4GB compressed) space. Finally, R250 is preprocessed in 14min, consuming only 0.7GB (0.17GB compressed) space. In general, CFLAT has an average preprocessing requirement of 3.306sec and 2.521MB per landmark.

Preprocessing Requirements @ Germany. The preprocessing requirements for the constructing the summaries of CFLAT in Germany, for various sizes of R-landmark sets, are shown in Figure 6. In general, there is a requirement for 29.33sec and 26.8MB per landmark, which is totally justifiable compared to Berlin, due to the larger size of the instance (by an order of magnitude). A significant improvement over the preprocessing requirements of FLAT is again achieved. E.g., for R2K FLAT requires (uncompressed) space 100.7GB which are constructed in 44.6h, whereas CFLAT creates the analogous preprocessed data in 16.3h requiring 53.6GB (16.1GB compressed) space. This indeed made it possible to consider landmark sets of size up to $4,000$ in the present work.

D Detailed Auditing of CFCA(N)'s Performance

We provide in this section more detailed experiments for the performance f CFCA. We start with mixtures of BC- and R-landmark sets. As Figure 7 shows, BC-landmarks improve mainly the relative error, whereas R-landmarks improve the query-time in Berlin. Interestingly, the best query-time is achieved by the hybrid landmark sets BC8K+R8K and BC4K+R12K, with the former having much better relative error.

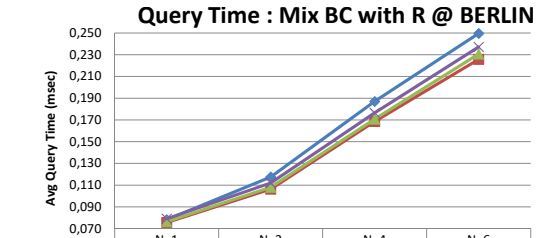

	N=1	N=2	N=4	N=6
BC16K	0,078	0,118	0,187	0,250
BC8K + R8K	0,076	0,106	0,169	0,226
BC4K + R12K	0,076	0,108	0,171	0,231
R16K	0,079	0,112	0,177	0,237

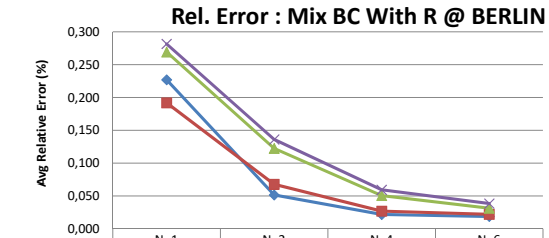

	N=1	N=2	N=4	N=6
BC16K	0,227	0,052	0,022	0,019
BC8K + R8K	0,192	0,068	0,027	0,022
BC4K + R12K	0,269	0,123	0,051	0,031
R16K	0,282	0,136	0,059	0,038

Figure 7 Performance of CFCA for mixtures (BC- and R-landmark types) of 16K landmarks in Berlin, and a query set of $50,000$ *random queries*.

Analogous observations hold also for Germany, as it is shown in Figure 8. Once more the best query-time (0.683msec) of CFCA is achieved for BC3K+R1K.

We next audit the amount of computational effort (both in terms of Dijkstra rank, and of absolute running times) of CFCA among its major steps.

Figures 9 and 10 give these measurements of CFCA(N) in Berlin. I.e., the number of arcs checked for relaxation by the initial TDD-ball from (o, t_o) in Step 1, the number of marked arcs connecting predecessors to intermediate vertices in Step 2, and the number of arcs checked for relaxation during the extension of the TDD-ball within the marked subgraph, in order to provide the resulting *od*-path.

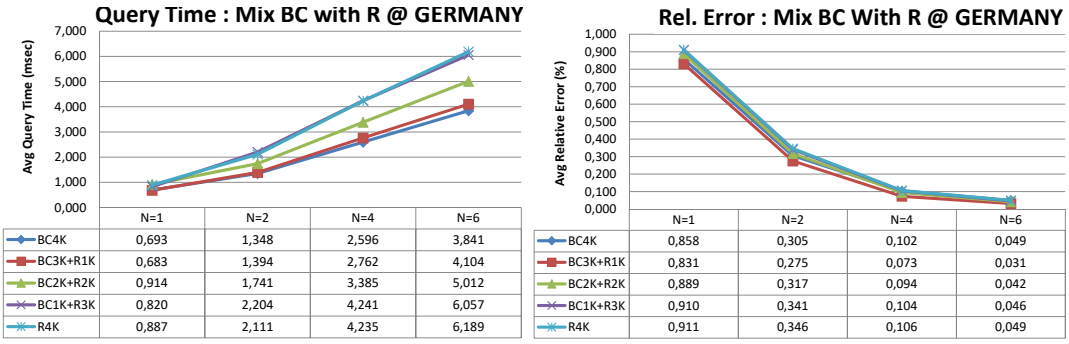

	N=1	N=2	N=4	N=6
BC4K	0,693	1,348	2,596	3,841
BC3K+R1K	0,683	1,394	2,762	4,104
BC2K+R2K	0,914	1,741	3,385	5,012
BC1K+R3K	0,820	2,204	4,241	6,057
R4K	0,887	2,111	4,235	6,189

	N=1	N=2	N=4	N=6
BC4K	0,858	0,305	0,102	0,049
BC3K+R1K	0,831	0,275	0,073	0,031
BC2K+R2K	0,889	0,317	0,094	0,042
BC1K+R3K	0,910	0,341	0,104	0,046
R4K	0,911	0,346	0,106	0,049

Figure 8 Performance of `CFCA` for mixtures (BC- and R-landmark types) of 4K landmarks in Germany, and a query set of 50,000 *random queries*.

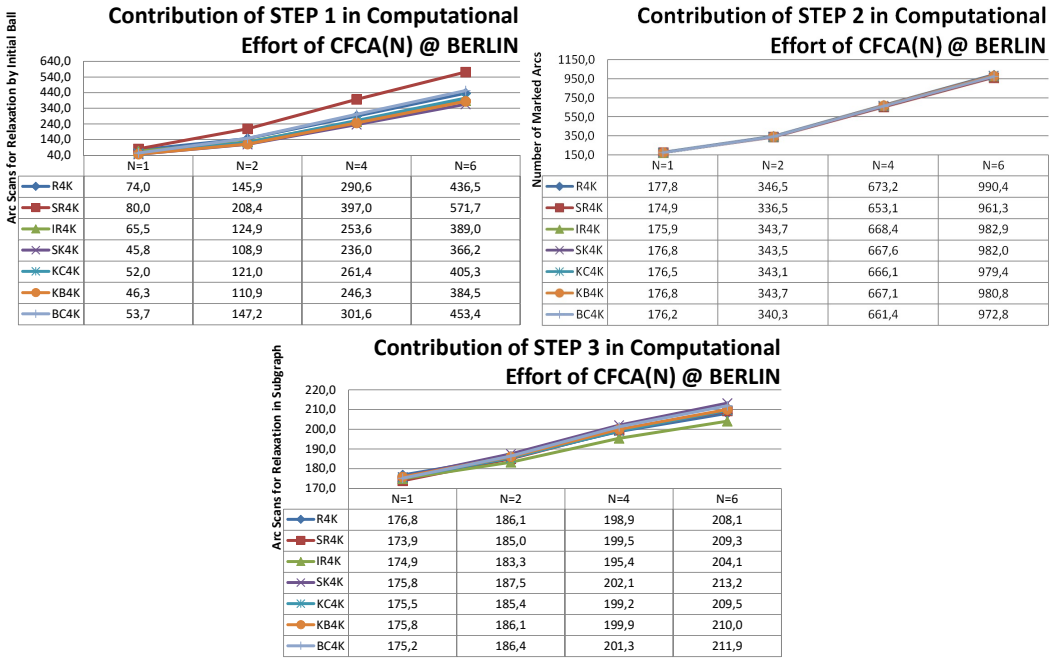

	N=1	N=2	N=4	N=6
R4K	74,0	145,9	290,6	436,5
SR4K	80,0	208,4	397,0	571,7
IR4K	65,5	124,9	253,6	389,0
SK4K	45,8	108,9	236,0	366,2
KC4K	52,0	121,0	261,4	405,3
KB4K	46,3	110,9	246,3	384,5
BC4K	53,7	147,2	301,6	453,4

	N=1	N=2	N=4	N=6
R4K	177,8	346,5	673,2	990,4
SR4K	174,9	336,5	653,1	961,3
IR4K	175,9	343,7	668,4	982,9
SK4K	176,8	343,5	667,6	982,0
KC4K	176,5	343,1	666,1	979,4
KB4K	176,8	343,7	667,1	980,8
BC4K	176,2	340,3	661,4	972,8

	N=1	N=2	N=4	N=6
R4K	176,8	186,1	198,9	208,1
SR4K	173,9	185,0	199,5	209,3
IR4K	174,9	183,3	195,4	204,1
SK4K	175,8	187,5	202,1	213,2
KC4K	175,5	185,4	199,2	209,5
KB4K	175,8	186,1	199,9	210,0
BC4K	175,2	186,4	201,3	211,9

Figure 9 Comparison of number of "touched vertices" per step of `CFCA(N)`, at 1.32sec resolution, for a query set of 50,000 *random queries* in Berlin.

It is clear from Figures 9 and 10 that only Step 1 depends on the type of landmarks that we consider. Observe also that Step 3 is essentially independent of the value of N, whereas the other two steps depend linearly on it. It is worth noting that, for R4K, while the contribution of Step 1 to the overall effort of `CFCA`, as N increases, varies from 17.3% to 26.7% w.r.t. the number of touched vertices, w.r.t. absolute times it is much more significant, varying from 35.6% up to 54.6%. This is exactly why we get a significant reduction in the query time when increasing the number of landmarks from 4K to 8K, but the (still significant) gain decreases as we go from 8K to 16K and almost vanishes when we go from 16K to 32K landmarks (cf. Figure 2). At least with respect to query-times, it seems that 16K is actually the ultimate size at which we should stop. On the other hand, the relative error keeps improving almost linearly with the number of landmarks.

Recall that the measurement does not only concern the estimation of an upper-bound on

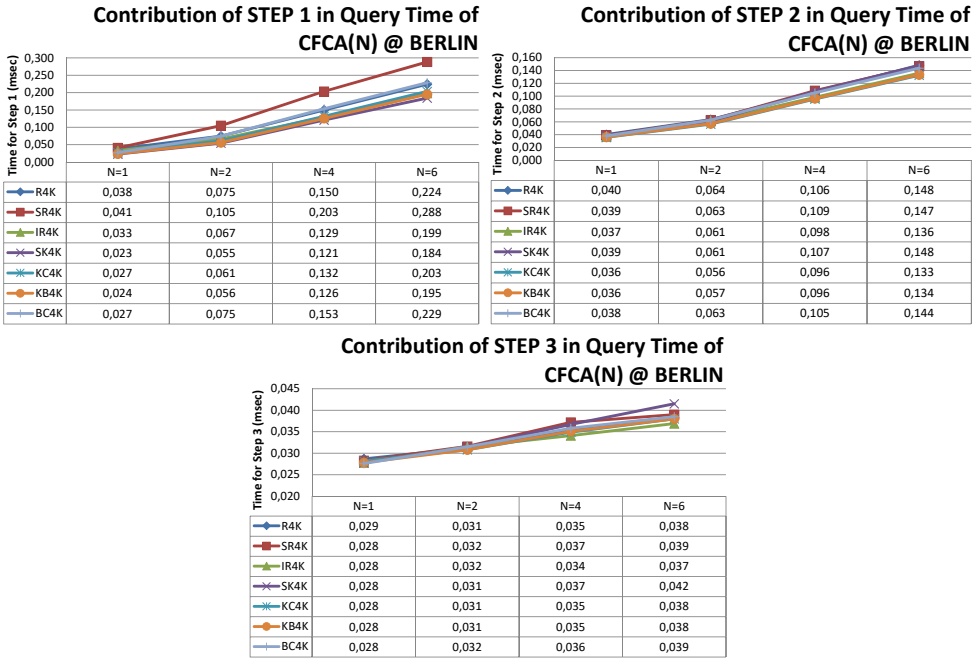

	N=1	N=2	N=4	N=6
R4K	0,038	0,075	0,150	0,224
SR4K	0,041	0,105	0,203	0,288
IR4K	0,033	0,067	0,129	0,199
SK4K	0,023	0,055	0,121	0,184
KC4K	0,027	0,061	0,132	0,203
KB4K	0,024	0,056	0,126	0,195
BC4K	0,027	0,075	0,153	0,229

	N=1	N=2	N=4	N=6
R4K	0,040	0,064	0,106	0,148
SR4K	0,039	0,063	0,109	0,147
IR4K	0,037	0,061	0,098	0,136
SK4K	0,039	0,061	0,107	0,148
KC4K	0,036	0,056	0,096	0,133
KB4K	0,036	0,057	0,096	0,134
BC4K	0,038	0,063	0,105	0,144

	N=1	N=2	N=4	N=6
R4K	0,029	0,031	0,035	0,038
SR4K	0,028	0,032	0,037	0,039
IR4K	0,028	0,032	0,034	0,037
SK4K	0,028	0,031	0,037	0,042
KC4K	0,028	0,031	0,035	0,038
KB4K	0,028	0,031	0,035	0,038
BC4K	0,028	0,032	0,036	0,039

Figure 10 Comparison of absolute running times per step of CFCA(N), at 1.32sec resolution, for a query set of 50,000 *random queries* in Berlin.

the earliest-arrival-time at (or equivalently, the shortest travel-time towards) the destination, but also the explicit construction of the corresponding *od*-path that guarantees this bound. Observe also that in absolute running times the speed-up is almost double, because the computationally most demanding step 2 only concerns accesses to the preprocessed data and there is no need for handling priority queues. Moreover, step 3 only concerns a very limited subgraph, containing only a few hundreds of arcs in overall.

Figure 11 demonstrates the analogous measurements for Germany. Again we observe the remarkable stability (and independence of the landmark set) for steps 2 and 3, as well as the linear dependence of steps 1 and 2, and the independence of step 3 on the value of N.

Observe finally that for Germany the speedups within the two measures (absolute running times, and "touched" arcs) are analogous. This is due to the fact that, since we have a quite small landmark set size this time, step 1 actually dominates the computational effort in this case.

E Exploring Outliers in Relative Errors

The purpose of our next experiment was to delve into the details of the relative error of CFCA(N). We study the quantiles of the relative error for serving 50,000 random queries, for BC16K at Berlin, and for BC4K at Germany. Figure 13 presents the results of this experimentation.

It is worth mentioning that in Berlin, with BC16K-landmarks we can have almost 99.52% of queries with error less than 1%, and 97.96% with error less than 0.1%. As for Germany, for BC4K-landmarks we can have 98.6% of the queries answered with an error less than 1% and 94.9% of them with error less than 0.1%.

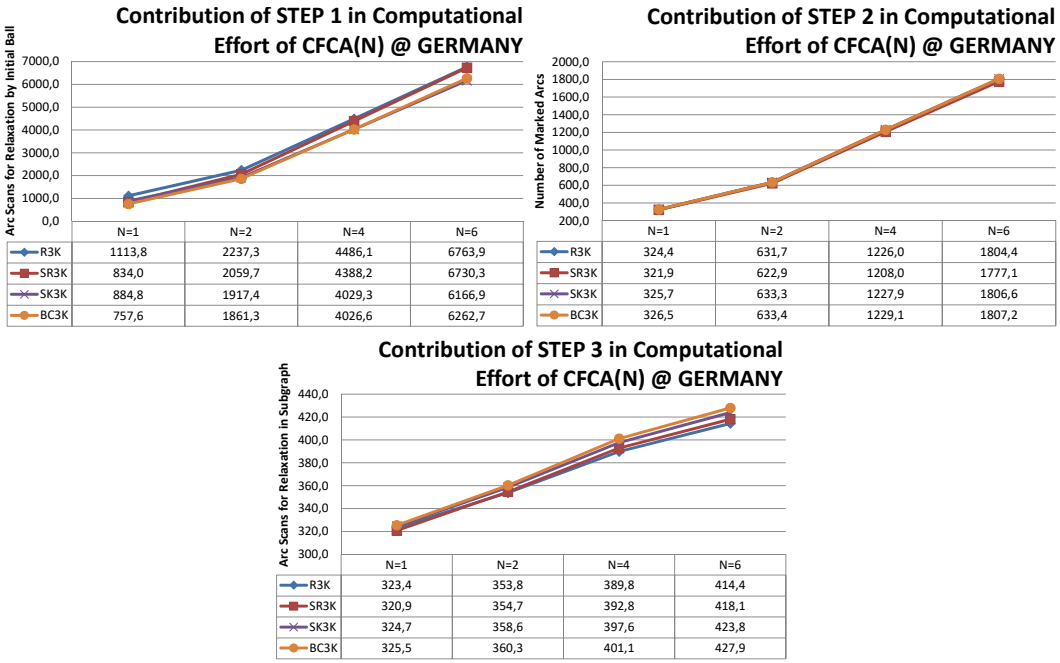

	N=1	N=2	N=4	N=6
R3K	1113,8	2237,3	4486,1	6763,9
SR3K	834,0	2059,7	4388,2	6730,3
SK3K	884,8	1917,4	4029,3	6166,9
BC3K	757,6	1861,3	4026,6	6262,7

	N=1	N=2	N=4	N=6
R3K	324,4	631,7	1226,0	1804,4
SR3K	321,9	622,9	1208,0	1777,1
SK3K	325,7	633,3	1227,9	1806,6
BC3K	326,5	633,4	1229,1	1807,2

	N=1	N=2	N=4	N=6
R3K	323,4	353,8	389,8	414,4
SR3K	320,9	354,7	392,8	418,1
SK3K	324,7	358,6	397,6	423,8
BC3K	325,5	360,3	401,1	427,9

Figure 11 Comparison of contributions in number of touched vertices, per step of CFCA(N), at 1.32sec resolution, for a query set of 50,000 *random queries* in Germany.

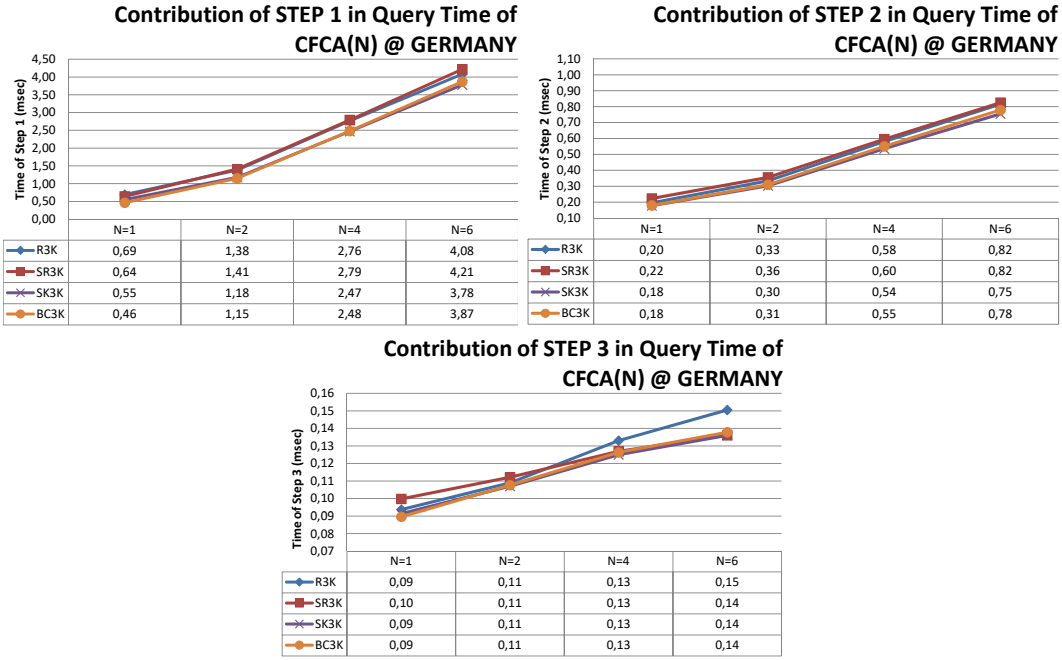

	N=1	N=2	N=4	N=6
R3K	0,69	1,38	2,76	4,08
SR3K	0,64	1,41	2,79	4,21
SK3K	0,55	1,18	2,47	3,78
BC3K	0,46	1,15	2,48	3,87

	N=1	N=2	N=4	N=6
R3K	0,20	0,33	0,58	0,82
SR3K	0,22	0,36	0,60	0,82
SK3K	0,18	0,30	0,54	0,75
BC3K	0,18	0,31	0,55	0,78

	N=1	N=2	N=4	N=6
R3K	0,09	0,11	0,13	0,15
SR3K	0,10	0,11	0,13	0,14
SK3K	0,09	0,11	0,13	0,14
BC3K	0,09	0,11	0,13	0,14

Figure 12 Comparison of running times per step of CFCA(N), at 1.32sec resolution, for a query set of 50,000 *random queries* in Germany.

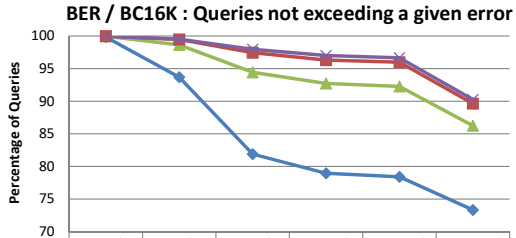

	< 10%	< 1%	< 0.1%	< 0.01%	< 0.001%	< 0.0001%
N=1	99,85	93,694	81,894	78,956	78,404	73,318
N=2	99,972	98,646	94,436	92,726	92,268	86,25
N=4	99,984	99,476	97,452	96,31	95,986	89,64
N=6	99,986	99,522	97,964	97,014	96,664	90,266

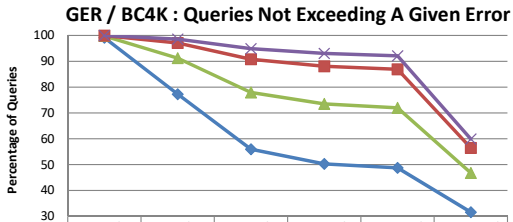

	< 10%	< 1%	< 0.1%	< 0.01%	< 0.001%	< 0.0001%
N=1	99,14	77,254	55,954	50,29	48,742	31,574
N=2	99,818	91,236	77,94	73,51	71,982	46,77
N=4	99,954	97,108	90,804	88,084	86,906	56,458
N=6	99,984	98,604	94,912	93,054	92,136	59,902

Figure 13 Tails of the error percentages of `CFCA(N)`, for 50,000 randomly chosen queries in the instance of Berlin with the BC16K landmark set, and for the instance of Germany with the BC4K landmark set.

F Discussion on State-Of-Art Oracles and Speedup Techniques.

`CFLAT` achieves a significant improvement compared to `FLAT` [15]. Concerning preprocessing requirements for Berlin (resp. Germany), `FLAT` consumed *compressed* space 14MiB (25.7MiB) and time 59.4sec (90sec), whereas `CFLAT` requires uncompressed space 2.58MB (27.44MiB) [or compressed space 0.702MiB (8.26MiB)] and 3.306sec (29.32sec), per landmark. As for the query performance, `FCA(1)` achieved 0.081msec (1.269msec) and 0.771% (1.534%), whereas `CFCA(1)` achieves 0.077msec (0.683msec) and 0.18% (0.831%), despite the fact that it also pays for the path construction.

We now proceed with the comparison of `CFLAT` with state-of-art speedup heuristics. In particular, we consider the speedup heuristics `inex.TCH` [3] (only for Germany), and `TDCRP` [4], `KaTCH`[5], `FreeFlow`, `TD-S` and `TD-S+A` from [22], `FLAT` [15], and `DijFreeFlow`, `CFLAT` in this work.

It should be once more noticed that `KaTCH`, `DijFreeFlow`, `FLAT` and `CFLAT` were experimented on our own machine, with exactly the same sets of uniformly and randomly selected queries. For the other algorithms we could only report (unscaled) the measurements of their experimentation by their authors, since we do not have the source codes at our disposal. For the sake of comparison and a posteriori verification, we provide the two random query sets that we have used in `http://150.140.143.218:8000/public/`.

For Berlin, the only experimentally evaluated speedup techniques we are aware of are `TDCRP` [4] and `TD-S`, `TD-S+A` [22]. We have also experimented with `KaTCH`, but the observed performance is dominated by most of the other algorithms, except for `TD-S+A`. `TDCRP` requires 21min of preprocessing time on a 16-core machine, 31MiB of preprocessing space, and achieves query performance (average query-time and relative error) 0.28msec and 1.47%. For an analogous amount of preprocessing *work*, `CFLAT` preprocesses R500 in 27min, exploiting 12 threads on a 6-core machine, consuming 1.3GiB (0.34GiB compressed) space. It achieves query performance varying from 0.356msec and 1.915% (for $N = 1$), to 1.848msec and 0.102% (for $N = 6$). If query-time is the main goal, then with BC8K+R8K `CFCA` achieves query performance varying from 0.076msec and 0.192% (for $N = 1$), to 0.226msec and 0.022% (cf. Figure 7).

The sweet spot of `CFLAT` w.r.t. the trade-off between query performance and preprocess-

[5] `https://github.com/GVeitBatz/KaTCH` with checksum 70b18ad0791a687c554fbfe9039edf79bc3a8ff3.

ing, seems to be for 4K landmarks: BC4K is preprocessed in 3h44min consuming 10.4GiB (2.8GiB compressed) space (cf. Figure 6) and CFCA then achieves query-performance varying from 0.088msec and 0.521%, to 0.367msec and 0.021% (cf. Figure 7). Moreover, for BC16K CFCA(6) provides stretch less than 1% for 99.522% of the 50,000 queries (cf. Figure 13). As for FreeFlow, TD-S and TD-S+A [22], it is certainly the case that these are quite simple algorithms which achieve remarkable performances. Their rationale is analogous to that of CFLAT: Certain paths for carefully selected time-windows (rather time-stamped shortest-path trees of CFLAT) are chosen, whose arcs induce a quite small subgraph in which TDD is executed. The difference with our oracle is that, instead of having the combinatorial structures automatically positioned in time, based on the time-dependent metric (as CFLAT does in order to achieve a required approximation guarantee), manual selection of time-windows (by trial-and-error) is used during the preprocessing of TD-S and TD-S+A. For running times, CFLAT can be faster than all these algorithms, e.g., for BC8K+R8K and CFCA(1). Concerning their remarkable error performances, it should be noted that for FreeFlow we tried to verify the reported errors by running our own version (DijFreeFlow). DijFreeFlow is not based on CH, but on running (static) Dijkstra on the free-flow metric and then computing the time-dependent travel-time along the chosen path. At least for the common query-set that we use in all our experiments, the error guarantees for FreeFlow are much worse than the ones reported in [22].

For Germany, we compare CFLAT with all the considered oracles and speedup heuristics. TDCRP requires total preprocessing time 4h41min on a 16-core machine, using 0.361GiB preprocessing space, and achieves query performance 1.17msec and 0.68%. inex.TCH(0.1), on the other hand, preprocesses the instance in 6h18min, consuming 1.34GiB space, and achieves query performance 0.7msec and 0.02%, and worst-case error 0.1%. For an analogous amount of preprocessing *work*, CFLAT preprocesses R1K in 8h9min using 12 threads of our 6-core machine consuming 26.8GiB (8.1GiB compressed) space, cf. Figure 6. CFCA achieves query performance varying from 2.175msec and 1.582% (for $N = 1$), to 11.974msec and 0.071% (for $N = 6$). If query-time is the main goal, then CFLAT preprocesses the hybrid landmark set BC3K+R1K in 32h35min consuming 107.2GiB (32.3GiB compressed) space, see Figure 6. CFCA achieves then query performance varying from 0.683msec and 0.831% (for $N = 1$), to 4.104msec and 0.031% (for $N = 6$), see Figure 8. Moreover, for BC4K CFCA(6) provides an error at most 1% for 98.604% of the 50,000 queries (cf. Figure 13). KaTCH is clearly worse than CFLAT. Indeed, the performance of KaTCH significantly deviates from the reported performances of all variants of inex.TCH, and is dominated by all oracles and speedup heuristics. One possible explanation might be that our own query set triggered some sort of bug in KaTCH, but it is impossible for us to verify this.

Finally, the query performances of FreeFlow, TD-S and TD-S+A for Germany are comparable to those of CFLAT, but the reported errors are much better. Again, we tried to verify the reported errors by running our own version (DijFreeFlow). At least for the common query-set that we use in all our experiments, the error guarantees for FreeFlow are much worse than the ones reported in [22].

Concerning temporal changes in the time-dependent data, the live-traffic updating procedure of CFLAT's preprocessed data, among 1,000 15-min randomly chosen disruptions, takes (per disruption) 0.275sec in Berlin for updating on average 48 affected BC4K-landmarks, and 37.676sec in Germany for updating on average 4 affected BC3K-landmarks (cf. Section G).

G Live Traffic Updating

As was done in [15], we conducted an experiment to assess the responsiveness of `CFLAT` to live-traffic updates. In particular, the goal is, when a disruption occurs "on the fly" (e.g., the abrupt and unforeseen congestion, or even blockage of a road segment for half an hour due to a car accident), how fast the oracle can take into account, for the affected route plans that have already been suggested or will be suggested in the near future, the temporal traffic-related information. We thus consider dynamic scenarios where there is a stream of live-traffic reports about abnormal delays on certain road segments (arcs), along with a time-window $[r_s, r_e]$, of typically small duration, in which the disruption occurs.

Our update step involves the recomputation of min-travel-time-path summaries for a subset of landmarks in the vicinity of the disruption. In particular, for a disrupted arc $a = uv$ of disruption duration $[r_s, r_e]$, we run a (static) `Backward-Dijkstra` from u under the free-flow metric, with travel time radius of at most $r_e - r_s$. The limited travel time radius is used to trace only the nearest landmarks that may actually be affected by the disruption, leaving unaffected all the "faraway" landmarks. The goal is to update as soon as possible the recommendations for the drivers who are close to the area of disruption. For each affected landmark ℓ, we consider a *disruption-times window* $[t_s, t_e]$, containing the latest departure-times from ℓ for arriving at the tail u at any time in the interval $[r_s, r_e]$ in which the disruption occurs. We then compute *temporal* travel-time summaries for each affected landmark and disruption-times window. This computation is conducted as in the preprocessing phase. Using a 15-min radius for the disruptions, we executed $1,000$ live-traffic updates for the instances of Berlin and Germany, for the landmark set BC4K and BC3K, respectively. For Berlin, the average number of affected landmarks was 48 for Berlin, and the updating procedure of the affected landmarks' summaries requires average time 0.275sec, using 12 threads on our 6-core machine. As for Germany, the average number of affected landmarks was only 4, and the updating procedure of the affected landmarks' summaries requires average time 37.68sec, again using 12 threads on our 6-core machine.